BASEMENT

MASSACHUSETTS
INSTITUTE of TECHNOLOGY
LIBRARIES

HD28
.M414
no. 1513-84
c. 2

DEWEY LIBRARY
LIBRARIES

WORKING PAPER
ALFRED P. SLOAN SCHOOL OF MANAGEMENT

THE WINNER'S CURSE IN BILATERAL NEGOTIATIONS

William F. Samuelson and Max H. Bazerman

September, 1984

Revised
WP 1513-84

MASSACHUSETTS
INSTITUTE OF TECHNOLOGY
50 MEMORIAL DRIVE
CAMBRIDGE, MASSACHUSETTS 02139

THE WINNER'S CURSE IN BILATERAL NEGOTIATIONS

William F. Samuelson and Max H. Bazerman

September, 1984

Revised
WP 1513-84

RESEARCH IN EXPERIMENTAL ECONOMICS:
A RESEARCH ANNUAL, JAI Press, Inc., in press

THE WINNER'S CURSE IN BILATERIAL NEGOTIATIONS

William F. Samuelson and Max H. Bazerman

Boston University

and Massachusetts Institute of Technology

1. Introduction.

The economic theory of bargaining is built on the cornerstone of self-interest: when the opportunity arises, profit-seeking individuals will negotiate and attain mutually beneficial agreements. Indeed, negotiation is often viewed as an analogue of or substitute for competitive markets — that is, under the right conditions bargaining will generate Pareto efficient economic allocations, as will perfectly competitive markets.[1] But "perfect" negotiations presuppose a number of conditions, in particular, that negotiators are perfectly rational and have perfect information about the bargaining situation. However, Bazerman and Neale (1983; Bazerman, 1983) have provided substantial evidence that negotiators deviate from rationality in systematically predictable ways. Short of the ideal of fully rational behavior, how will negotiations proceed? How should an individual negotiate when only imperfect or limited information about the negotiation setting is available? When he or she has worse information than the other side and both know it? What negotiating procedures are successful in reaching mutually beneficial agreements?

This paper addresses these questions and presents experimental evidence on bilateral bargaining behavior under uncertainty. A main finding is that under asymmetric information, negotiators systematically deviate from normative behavior and, consequently, fall prey to the "winner's curse" — that is, they consistently (and voluntarily) enter into loss-making purchases and forego profit-making opportunities. These losses result because subjects act as if the other party has the same information as themselves. For example, many uninformed buyers make expected losses on transactions because they fail to anticipate the informed opponent's profit-maximizing bargaining behavior. In

turn, informed sellers, when given the opportunity to make price offers, fail to take full advantage of their information advantage. In short, both negotiators — informed and uninformed party alike — fail to recognize the force of the information asymmetry.

The present study extends the analysis of the winner's curse to bilateral negotiations. This phenomenon has been examined for some time in the area of competitive bidding. Researchers (e.g., Capen, Clapp, and Campbell, 1971; Bazerman and Samuelson, 1983) have asserted that it is common for "winning" bidders in competitive auctions to find that they have overpaid for the acquired commodities. The key point is that bidders tend to ignore the impact of competitor bidding behavior on their own optimal strategy.

Additional research supports the more general conclusion that competitive decision makers systematically ignore the impact that the decisions of other parties can have with respect to their optimal behavior. A good example is the dollar auction exercise (Shubik, 1971; Teger, 1980) in which the highest bidder in an oral auction pays its bid and receives the dollar, and the second highest bidder pays its bid and receives nothing in return. The common result in this auction is an escalating pattern in which individuals bid far in excess of a dollar and which produces significant profits for the auctioneer. Why do bidders get involved? One explanation is that individuals see the potential for profit early in the auction, and fail to take the perspective of what the auction will look like to other bidders. If the bidder considers the dollar auction from the point of view of both competitors, it is easy to see the benefit of staying out of the auction. Overall, there is substantial evidence that competing decision makers fail to consider the impact of the conditional behavior of other parties in making decisions.[2]

The analysis of the winner's curse in the negotiation setting relies on the follow example:

> One firm (the acquirer) is considering making an offer to buy out another firm (the target). The complication is that the acquirer is uncertain about the ultimate value of the firm. Though it has reason to believe that the target will be worth more under its management than under present ownership, the acquirer (even after making its best estimates) does not know the target's ultimate value. Target management, on the other hand, has an accurate estimate of value and so shares none of the acquirer's uncertainty. In these circumstances, what final price offer should the acquirer make for the target? Alternatively, if the target company has the opportunity to make a take-it-or-leave-it offer, what price should it name? Should the acquirer accept it?

Two features are present in the example above. First, the aspect of the problem of greatest interest to economists is the opportunity for mutual gain. The fact that the company is more valuable in the hands of the acquirer than under present management means that ex post there is a price at which both sides can profit from the sale. Second, the value of the transaction is uncertain, and one side has different (better or worse) information about the uncertain value than the other. For example, the target firm's management ordinarily has proprietary information about its operations which is not available via published accounting reports or in other forms to the acquirer or other outsiders. (Of course, the acquirer may also have information about the firm's value under new management which the target doesn't have.)

This asymmetry of information should put the less well-informed party on guard. For instance, the management board of the acquiring firm might reason as follows:

> The first fact is that we are uncertain about the ultimate value of the target. If we extend an offer to target management and it is accepted, are we going to be sold a 'bill of goods'? After all, it is more likely that target management will unload an ailing company than a healthy one. By making a low offer, will we simply be getting what we paid for or even less than what we paid for? What constitutes a profit-maximizing offer? Indeed, is there any offer which, if accepted, will provide a positive expected profit from the transaction?

In the reasoning above, the acquirer recognizes the presence of asymmetric information and anticipates the behavior of the better-informed target. Though this reasoning is easy to follow, the evidence suggests that individuals consistently fail to develop this logic when faced with competitive situations. Rather, an alternative model is proposed to describe the behavior of a significant portion of competitive decision makers. In this model, negotiators behave as if the opponent possesses the same information as themselves -- even when they are told that the opponent has better or worse information available. Negotiators make this "naive" assumption about the information and behavior of the opponent in order to simplify their decisions in a complex environment.

This paper presents the results of a series of bargaining experiments designed to test equilibrium and naive models of bargaining behavior under uncertainty, where, as in the example above, one party is less well-informed than the other. Our approach and conclusions are both normative and positive in nature. The normative analysis offers a prescription for a rational individual's optimal bargaining strategy and examines the welfare implications of such behavior. As will be shown, even very simple bargains under uncertainty require individuals to make subtle probabilistic inferences about the potential value of the transaction. Moreover, in contrast to agreements under certainty, bargains under uncertainty present a direct conflict between individual self-interest and group welfare. Optimal bargaining behavior may preclude the attainment of mutually beneficial agreements even when it is common knowledge that such agreements exist. In short, a normative analysis identifies imperfect information as a potential barrier to mutually beneficial transactions.

A positive analysis investigates the actual bargaining behavior of individuals under laboratory conditions. The experiments employed a pair of bargaining procedures, differing as to whether the uninformed buyer or the informed seller initiates the price offer. The results show that under either bargaining method, subjects deviate from normative guidelines in predictable ways. For instance, in formulating price offers, buyers (acquirers) consistently fail to appreciate their information disadvantage, and thereby sacrifice profits. The more difficult the inferential task, the greater are bargainers' deviations from appropriate normative behavior. Indeed, in one bargaining experiment, the vast majority of subjects actually pursued bargaining strategies predicted to generate expected losses on average -- a strategy which no rational (risk neutral or risk averse) bargainer would voluntarily pursue. Similarly, informed sellers, when given the opportunity to make price offers, failed to take full advantage of their information advantage.

A final result is that actual subject bargaining behavior, though individually suboptimal, is collectively advantageous. Because observed bargaining strategies are more "cooperative" than the normative benchmark, the frequency of agreements exceeds that which would pertain in equilibrium. In particular, the expected group profit is maximized when the informed party has the opportunity to make a price offer.

The paper is organized as follows: section two introduces and discusses three different versions of the basic negotiation model used throughout the paper. Section three provides normative and "naive" analyses of each version. Section four presents and analyzes subject behavior in an experimental test of the negotiation model. Finally, section five provides some concluding remarks and discusses the applicability of the laboratory findings.

2. The Negotiation Model.

The basis for the experiments that were conducted consists of three versions of a short bargaining exercise entitled "Acquiring a Company", which are reproduced in the appendix. The generic elements of the bargaining model underlying the exercise can be formally described as follows:

> A target and potential acquirer are negotiating over the sale of a good of uncertain value. Denote the monetary value of the good to the target by v. This value is known by the target but not by the acquirer who regards v as a random variable with cumulative probability distribution F(v). In turn, denote the value of the good to the acquirer by w(v). This functional notation indicates that the value of the good to the target may vary with the value to the acquirer. Whatever the value of v, the good is always worth at least as much to the acquirer than to the target-- that is, $w(v) \geq v$ for all possible v. Both sides know the functions F(v) and w(v), but only the target knows the specific values of v and w. Concerning these last two values, the acquirer has only the probabilistic information given by F(v).

In the experiments, three different versions of the negotiation exercise were used. These were created by varying slightly the parameters of the negotiation situation. Employing the notation introduced above, the main features of each of the versions can be summarized as follows:

Version 1. $w(30) = 30$, $w(60) = 130$ and $F(v) = \begin{cases} 0 & \text{for } v < 30 \\ 1/2 & \text{for } 30 \leq v < 60 \\ 1 & \text{for } v \geq 60 \end{cases}$

Version 2. $w(v) = v + 30$ and $F(v) = v/100$, for $v \in [0,100]$.

Version 3. $w(v) = 1.5v$ and $F(v) = v/100$, for $v \in [0,100]$.

Version 1 is the simplest of the three and employs a discrete probability distribution over two values of v. Note that the acquirer's value is strictly greater than the target's value only in the case that $v = 60$. In turn,

versions 2 and 3 share a uniform distribution of possible values but differ in the functional form describing the acquirer's "absolute advantage" for the good, w(v) - v. In Version 2, the acquirer's advantage is constant for all v. In Version 3, the advantage is proportional to v.[3]

The bargaining procedures employed in the experiments are of a very simple kind. One party, either the acquirer or the target, makes a "first and final" price offer which the other can accept or reject.[4] If the offer is accepted, a sale takes place at the offer price. If not, there is no sale, and no money changes hands. In corporate acquisitions, the most common practice is for the acquiring company to make a tender offer which the target can accept or reject. Alternatively, the target management, cognizant of the firm's true value, could name a "buy out" price. In the experiments, both procedures, acquirer and target offers, were used.

3. Normative and Naive Hypotheses.

The analysis of the experimental results relies on two competing hypotheses about subject bargaining behavior. Under normative behavior, subjects correctly account for the presence of information asymmetry and employ optimal (i.e., equilibrium) bargaining strategies. Under naive behavior, bargainers employ simpler strategies which ignore the information asymmetry. Specifically, the naive model predicts that bargainers will act as if the opponent has the same information as themselves. We examine both normative and naive hypotheses under 1) acquirer "bids" and 2) target offers.

Acquirer Bids (Normative Behavior). First, consider the acquiring company's choice of bid. What price should it name for the target company's shares in each of the three versions of the bargaining model? Given the facts

of Version 1, the management board of the acquiring company should reason as follows:

> Any price bid between \$30 and \$60 will be accepted only by a low valued company worth \$30. In this range, no profit is possible and the higher the price bid the larger is the possible loss. However, a price in excess of \$60 per share will <u>always</u> be accepted by the target. Since the average acquisition value is (1/2)(30) + (1/2)(130) = \$80 per share, a profit is possible at any price in the range \$60 to \$80 per share. Clearly, the most profitable bid is \$60 per share (or if necessary, \$60 + ε where ε is in the positive neighborhood of zero).

Thus, in evaluating the behavior of subjects, the range of bids can be partitioned into three intervals -- \$0 to \$59 per share, \$60 to \$80 per share, and above \$80 per share. Only for bids in the middle interval does the acquirer earn a non-negative expected profit, and at the lower limit of this interval, \$60 per share, it earns the maximum expected profit, \$20 per share.

The chain of reasoning above depends on three presumptions: 1) The target acts to maximize profit and therefore accepts the acquirer's price P_A if and only if $P_A \geq v$; 2) The acquirer anticipates the target's acceptance behavior and correctly revises its probability assessment of the company's acquisition value <u>conditional on its bid being accepted</u>; and 3) The acquirer prefers to bid \$60 per share (and face the equally likely profits of 130 - 60 = \$70 per share <u>or</u> 30 - 60 = -\$30 per share) rather than bid \$0 per share. Though the simple logic of each of these steps is obvious, recognition of this line of reasoning is more subtle.

Similar reasoning can be applied to determine the acquirer's optimal bid in Version 2. The management board reasons as follows:

> Suppose we make a bid of (let us say) $30 per share. How often will this bid be accepted? The answer must be thirty percent of the time, since the target will accept if and only if the value under current management v is less than the bid. (All values between $0 and $100 per share are equally likely). What is the average value under current management of companies thus sold? Fifteen dollars per share, since all values between $0 and $30 are equally likely. The average value under our (the acquirer's) management? Forty-five dollars per share, since the company is worth $30 per share more in our hands. Thus, in the event of an agreement, the average profit is $45 - $30 = $15. In turn, the acquirer's overall expected profit is (.3)(15) or $4.5 per share.

Once again, profitable bargains are possible. The acquirer's expected profit is zero at a price of $0 per share and also at $60 per share (as is easily checked). For bids between these values, the firm earns a positive expected profit, and above a price of $60, it's expected profit is negative. The acquirer's maximum expected profit occurs at P_A = $30 per share (halfway between the "break-even" prices). In fact, the bid P_A = $30 stochastically dominates any other bid -- that is, it offers strictly better odds of better profit outcomes.[5] Thus, any profit maximizer, regardless of his risk attitude, should choose P_A = $30.

There is a subtle but significant difference between the second and third versions of the bargaining experiment. In Version 3, the company, whatever the target's value, is worth 50% per share more to the acquirer than to the target, i.e., w(v) = 1.5v, for all v. With this modification and employing a line of reasoning analogous to that of Version 2, the acquiring company's management board reasons as follows.

> A bid of (let's say) $60 per share will be accepted 60% of the time by targets with an average value of $30 per share. Thus, the average acquisition value of such a company is 50% more or $45 per share. If accepted, our profit from this bid is, thus, $45 - $60 or -$15 per share. Consequently, a $60 per share bid is ill-considered.

It is not hard to see that the same kind of reasoning applies to any positive price bid the acquirer might consider making. On average, the acquirer obtains a company worth 25% less than the price it pays. Thus, the acquirer's best bid is $0 per share, which, of course, is tantamount to making no bid.

Version 3 offers a graphic illustration of the tension between the opportunity for mutual gain and the impact of asymmetric information in simple bargaining situations. Even though in all circumstances the firm is worth 50% more to the acquirer than to the target, any tender offer that the acquiring firm might make results in a loss on average. Indeed, the source of this barrier to trade stems from the presence of adverse selection. A given bid will be accepted only by "low-value" companies with the result that the average value of acquisitions falls short of the purchase price.

To sum up, the uninformed acquirer's optimal price bids are $60, $30, and $0 in the respective versions of the bargaining experiment.

Acquirer Bids (Naive Behavior). In contrast, to the normative reasoning above, a "naive" negotiation strategy fails to account for the information asymmetry present in the bargaining setting. Specifically, a naive acquirer formulates its bid strategy as if the other side had the same information as itself. A naive acquirer reasons as follows:

> The value of the firm is uncertain and differs depending upon which management controls it. Given this uncertainty, the greatest price that we can pay is E[w(v)]. In turn, the lowest price that target management is willing to accept is E[v], (where it just breaks even on average). Thus, our optimal bid is $P_A = E[v]$.

The naive analysis relies on the unconditional expected values E[v] and E[w(v)], correctly identifying the latter value as the acquirer's maximum profitable purchase price but incorrectly perceiving the former as the target's minimum price. This last presumption would be correct only if the target, as well as the acquirer, were uninformed about the company's intrinsic

value. By failing to anticipate the informed target's acceptence behavior, a naive acquirer falls prey to the winner's curse making expected losses on those transactions completed.

In Version 1, a naive acquirer makes a price bid of $45 per share (the average value under target management) and suffers an expected loss of (.5)(45-30) = $7.5 per share since only "$30 companies" accept this bid. In Versions 2 and 3, a naive acquirer identifies $50 per share (the average value under target management) as its most profitable bid. This price, however, leads to a substantial loss in comparison to normative behavior. In Version 2, a $50 per share bid results in a (reduced) expected profit of (.5)($55-$50) = $2.5 per share; in Version 3, the acquirer makes an expected loss of (.5)($37.5-$50) = $6.25 per share. In each case, the naive strategy would be optimal if the target were also uninformed. For easy reference, Figures 2.1-2.3 show the central tendencies of price bids under normative and naive behavior as well as the acquirer's expected profit as a function of its price bid for the respective versions.

Target Offers (Equilibrium Analysis). A second part of the negotiation exercise reversed the roles of the experimental subjects allowing the informed target to name a first-and-final price which the acquiring company could accept or reject. The information available to the target and acquirer remained the same.

When the target names a first-and-final offer, the negotiation strategies of the players are interdependent — that is, one side's "best" strategy depends on the other side's strategy. In general, the informed target will name a price in excess of its true value. Clearly, the acquiring company's acceptance strategy is influenced by the target's pricing strategy.

Against an "aggressive" target (one who marks up its price well in excess of its true value), the acquiring company should use a strategy with a low rate of acceptance. In turn, the target's strategy depends on the buyer's acceptance behavior. Obviously, the target can push for higher prices if these are likely to be accepted.

Given this strategy dependence, our initial analysis focuses on Nash equilibrium strategies as the appropriate normative benchmark. A pair of player strategies is in equilibrium if neither side can unilateraly deviate from its chosen strategy and increase its expected payoff. Then, it can be shown that the target's optimal strategy is to name the same price P_T for all $v \leq P_T$.[6] In turn, the acquirer's optimal strategy is to accept any price up to and including P_T, but no higher price. Finally, the target determines the price P_T so that the acquirer just breaks even (earns a zero expected profit) by accepting it.

The key point to note about this equilibrium is that the target's price is not contingent on the value of v -- that is, the same price is named for all values of v within a specified range. The explanation for this result stems from the acquirer's acceptence behavior. The acquirer's natural strategy is to accept the target's demands within a pre-determined lower range -- to accept prices up to some value and reject any higher price. Since the acquirer is uninformed at the time it accepts or rejects the offer price, the target is free to name P_T for all $v \leq P_T$. (It is unprofitable to attempt to name any higher price since it won't ever be accepted.) In short, the acquirer knows and the target knows that naming the price P_T (while excluding cases such that $v > P_T$) does not affect the relative likelihood of any $v \leq P_T$. The target's problem is simply to determine the greatest price that the acquirer will accept.

In Version 1, the target's optimal price is $80 (simply the acquirer's expected value for the company). A risk-neutral acquirer finds this price barely acceptable and would make an expected loss for any higher price. In Version 2, the target's maximum price is $60. An acquiring company that accepts this price obtains companies worth between $0 and $60 under current management or between $30 and $90 to itself. Consequently, the acquirer obtains companies worth $60 per share on average and just breaks even on its acquisitions. Finally, in Version 3 any price that the target might offer will result in an expected loss for the acquirer. As in the case where the acquirer makes an offer, no agreements are possible in equilibrium.

To sum up, facing a risk-neutral acquirer, the target's optimal noncontingent price is $80 in Version 1 and $60 (for $v \leq \$60$) in Version 2. In Version 3, the target earns a $0 profit for all prices, since all prices will be rejected in equilibrium.

Target Offers (Naive Behavior). Given the chance to name a price, a naive target can be expected to reason as follows: The company's current value and potential acquisition values are v and w(v) respectively. Therefore, a mutually beneficial agreement is possible only at a price P_T such that $p_T \in [v, w(v)]$. In short, the hypothesis of naive behavior predicts that actual target offers will be bounded by these respective values.

Target Offers (Empirically-Best Strategies). Thus, far, we have described equilibrium and naive negotiation strategies. The strong appeal of equilibrium strategies is that each is a best response against the other side's best strategy. Under acquirer price bids, the equilibrium benchmark is particularly compelling. The target's dominant strategy is to accept any price not smaller than v. Thus, any rational acquirer should anticipate this

response and make an optimal bid accordingly. By contrast, under target company offers, neither side has a dominant strategy; each player's optimal bargaining strategy depends directly on the strategy employed by the other. Though equilibrium strategies continue to provide a normative benchmark, it is important to point out that better responses may be available against players who deviate from their equilibrium strategy. Thus, a negotiator can take advantage of an opponent's predicted deviations from the equilibrium strategy by employing an (empirically) best strategy in response. Raiffa (1982) terms this the asymmetrically prescriptive/descriptive approach, while arguing that the researcher should be interested in prescribing the behavior of a focal negotiator in view of the predicted actual behavior of the competing party. Thus, under target company offers, the target can push for higher prices if the acquirer is willing to accept prices in excess of the equilibrium price P_T. With this in mind, the examination of the experimental results not only compares actual player behavior to the equilibrium benchmark, but also considers each player's optimal strategy in response to the actual behavior of the population of experimental opponents.

4. Experimental Results.

The subjects in the experiments were 595 MBA students at Boston University (drawn from 18 Managerial Economics classes). Besides its research intent, the experiments served as an introduction to decision making under uncertainty. The exercises took from five to ten minutes for subjects to complete, after which the experimenters conducted a brief discussion. The bulk of the experiments did not employ monetary payoffs. This and the fact that the experiments were conducted in class permitted the largest possible data sample. At the same time, it was important to insure that subjects had

sufficient incentives to make serious decisions -- that is, decisions of the kind they would make were monetary returns at stake. To test subject incentives, a sub-sample of students participated in Version 3 of the exercise where monetary payoffs were at stake. Subject behavior under this experimental treatment (reported at the end of this section) was found to be statistically indistinguishable from behavior in the experiments without monetary payoffs. We begin with the results of the non-monetary experiments.

Acquirer Bids. The distributions of price offers for Version 1, 2, and 3 of the experiment are shown in Figures 1.1, 1.2, & 1.3 respectively. (In all versions, subjects overwhelmingly chose round numbers for bids. For instance, in Version 1, excluding the bid of $61, only 6% of the offers were not divisible by 5.)

First, consider the bids made in the simple discrete case of Version 1. Virtually all subjects viewed the acquisition opportunity as profitable, making "serious" bids, i.e., bids greater or equal to $30 which were potentially acceptable to the target. Normative behavior -- that is, bids in the interval $60-$65 inclusive -- was adopted by 41% of the subjects. At the same time, significant systematic deviations occured. As the figure indicates, the bid distribution is slightly bimodal -- the second most frequent price interval being $40-$45 per share. Indeed, nearly as many subjects made expected losses (from bids in the intervals $31-$59 and $81-$130) as achieved maximum profits (offers $60-$65).

In Versions 2 and 3, subjects adopted normative behavior even less frequently. In Version 2, an optimal (and stochastically dominant) price bid in the interval $30-$35 was submitted by only 16% of the subjects. Moreover, nearly 50% of the subjects made bids in excess of $60 per share, thereby incurring expected losses on their transactions. In Version 3, since all

positive offers are loss making, an optimal bid is $0. But the startling conclusion of Figure 1.3 is that the overwhelming majority of the subjects (some 93%) rejected normative behavior. Only 9 of 123 subjects submitted zero bids. Indeed, there were very few low bids of any kind. Apparently, the subjects perceived the bid opportunity as profitable and spurned low bids.

The evidence in all three versions of the experiments strongly supports the hypothesis of naive behavior on the part of a significant number of subjects. In Version 1, the naive bid, $45 per share, was the second most frequent price (after $60) and was submitted by some 16% of all subjects. In Versions 2 and 3, similar patterns emerge. In each case, substantial numbers of subjects submitted bids equal to the unconditional expected value of the company. Thus, in both versions, the most frequent category of bids was the interval $50-$55 per share (accounting for 22% of all bids in Version 2 and 35% in Version 3). Apparently, many subjects viewed $50 per share, the expected value of the firm under current management, as the lowest bid acceptable to the target and judged it, therefore, to be the most profitable offer. Further, 68% of the subjects in Version 2 made bids between the naive $50 per share price and the unconditional expected value of the company to the acquirer -- $80 per share. Similarly, 73% of the subjects in Version 3 made bids between the naive $50 per share price and the unconditional expected value of the company to the acquirer -- $75 per share. Thus, there is substantial evidence that many subjects looked for mutually profitable opportunities based on a naive model of negotiator behavior -- leading most of them to surprising losses.

To sum up, the general pattern of bids by uninformed buyers represented a mixture of normative behavior and naive behavior -- the latter being more frequent than the former except in Version 1. The proportions of subjects

exhibiting normative and naive behavior were both highly statistically significant. For instance, on the basis of a chi-square test, one would reject at the .01 significance level the hypothesis that the observed bids of Versions 1 and 2 were drawn randomly from the interval $30 to $80 per share. Similarly, the hypothesis of random bids between $0 and $75 would be rejected at the .01 level in Version 3.

Target Offers. The second section of the exercises asked subjects to report the prices they would offer acting on behalf of target management for alternative values of v and, in turn, the maximum price they would accept representing the acquiring company. In the target role, subjects reported prices for the possible values $30 and $60 in Version 1 and for the "representative" values $0, $20, $40, $60, $80, and $100 in Versions 2 and 3. Acting for the acquiring company, they reported the maximum price that they would accept -- that is, if the acquirer reported a maximum price $x, it would accept any target price smaller or equal to $x but no higher price.

Recall that a normative analysis predicts that the target should name a price which is invariant with respect to it's true value. In turn, the acquirer should name this same value as its maximum price (just breaking even on the acquisitions it accepts). In all versions, the evidence indicated that subjects deviated sharply from normative behavior in systematic directions. With few exceptions, targets displayed monotonic offer behavior -- that is, the target's price offer increased with the underlying value v. Thus, in contrast to the normative benchmark, target offers tended to transmit information about the true value of the company (the greater the offer the greater the company value). Moreover, the typical target's monotonic offer strategy, while reserving the lion's share of the profit for itself, left a small profit for the acquiring company.

Figure 3.1 depicts the distribution of target price offers for each of the values v = 30 and v = 60, as well as the distribution of acquirer acceptence prices for Version 1 of the bargaining exercises. Scrutiny of the prices offered by targets indicates that subjects overwhelmingly rejected equilibrium behavior (a constant price offer of $80 for high and low value companies alike). Instead, offers (though widely spread) depend closely on values. For v = 30, the mean price offer among subjects is $67.2; for v = 60, the mean offer is $97.3. Indeed, price offers are highly informative of values. Obviously, any offer smaller than $60 marks the target as having a value of $30 with certainty. At the same time, for very high demands, the relative likelihood is that the target has a high value (i.e., an acquisition value of $130). This is consistent with the prediction of naive target behavior in which target offers increase with values -- as would be optimal in the case that <u>both</u> parties were fully informed.

Following Raiffa's asymmetrically prescriptive/descriptive approach, let us consider the target's expected profit maximizing offers given the distribution of acquirer acceptence prices (see Figure 3.1b).[8] For a given v, the optimal offer is determined by trading off the probability and profitability of an agreement -- where the latter is increasing and the former is decreasing in the target's price. When v = 30, one calculates (through trial of all possible prices) that the optimal offer is $60 per share. At this price the target's expected profit is:

$$(60-30)\ \text{Prob(price is accepted)}$$

$$= (60-30)(.71) = \$21.3 \text{ per share,}$$

since 71% of the population of acquiring companies reported an acceptance price of $60 or higher (see Figure 3.1b). In turn, for v = 60 the target's optimal price offer (through trial of all possible prices) is $80 per share

(accepted 40% of the time) which generates an expected profit of (.4)(80-60) = $8.00 per share. (When holding v = 30, a signicant number of subjects made offers in the neighborhood of $60, the empirically best price; for v = 60, very few subjects adopted the best price of $80.)

Similar analyses pertain to the latter two versions of the bargaining exercise. Figures 3.2 and 3.3 display the price offers submitted by targets and the acceptence prices of acquirers for the respective versions. The figures provide strong evidence for the naive bargaining model. In each version, the vast majority of targets made offers in the interval [v, w(v)] -- consistent with the naive model. Indeed, one-half of all targets in Version 2 and one-third of the targets in Version 3 employed the offer strategy $P_T(v) = w(v)$, i.e., named a price equal to the company's value to the acquirer. (Note the accumulation of offers along the line $P_T = v+30$ in Version 2 and $P_T = 1.5v$ in Version 3.) In both versions roughly one-third of the price demands fell below w(v), thereby allowing the acquirer a profit on the transaction. This naive behavior would be appropriate if the acquirer had the same information as the target -- which the targets knew was not the case. In short, a naive target made price offers as if the acquirer had identical information as itself. The remaining prices were well in access of acquisition values, $P_T(v) > w(v)$. These prices occured mainly at low target values: v=0, v=20, v=40, and v=60. For instance, several subjects in each version made price offers of $80 and $100 per share for each of the values, 0, 20, 40, and 60 (before increasing their offers at higher values). To sum up, while most subjects employed naive (and monotonically increasing) offer strategies, a minority of subjects followed more aggressive strategies, such that $P_T(v) > w(v)$.

For easy reference, Tables 1.2 and 1.3 present actual mean price offers, equilibrium offers, and empirically-best offers (obtained by trial of all possible prices) of targets for Versions 2 and 3 respectively. For each value v, the empirically-best offer provides the target the maximum expected profit against the actual distribution of acquirer acceptence prices. In each version, empirically-best offers are increasing in values (though not strictly so). In Version 2, targets face a more attractive distribution of acceptence prices and so make higher optimal price offers and earn greater expected profits.

Experiments with Monetary Payoffs. As outlined in Smith (1982), the use of monetary payoffs has been the most frequent means of inducing prescribed values for outcomes on the part of subjects. Since the bulk of the present experiments did not employ monetary payoffs, it was important to test whether the type of reward structure — monetary or non-monetary — was crucial to the experimental behavior observed. Toward this end, a series of experiments were conducted employing monetary payoffs. These were identical to the non-monetary experiments with two exceptions: 1) each subject was assigned to a single role, acquirer or target, and 2) the instructions received by subjects included a section outlining how monetary payoffs would be determined. In addition, since monetary gains and (possibly) losses were at stake, participation in the exercise was purely voluntary. (In all, 19 of 131 subjects chose not to participate.) At the conclusion of the exercise (after all papers were collected), the experimenters employed simple randomizing devices (drawings from jars) to determine the bargaining outcome and the monetary payoffs to subjects. Monetary payoffs were at the rate of 10¢ per unit of profit earned. For instance, a bid of $50 per share by an acquirer (if accepted)

implies a maximum profit of (75-50) or $2.50 and a maximum loss of (0-50) or $5.00. The magnitude of reward was deemed sufficient given the amount of time (five to ten minutes for the typical subject) required of the decision task. The instructions for the monetary experiments are outlined in the appendix.

The distribution of price bids by acquirers under monetary payoffs are presented in Figure 4.3a. A comparison with Figure 1.3 indicates strongly that bids follow similar distributions whether or not monetary payoffs are present. In each case, the overwhelming majority of subjects (93% and 92%) submitted non-zero bids and, therefore, fell prey to the winner's curse. In each case, the most frequent offer is $50 per share (appropriate if both parties were uninformed of the company's true value) and the majority of bids lie between $50 and $75 per share. The sole difference between the distributions is that under monetary payoffs bids between $20 and $45 per share occur with greater frequency (at the expense of bids between $50 and $75) than under non-monetary payoffs Thus, there is the indication that when monetary payoffs are at stake, subjects tend to moderate their bids slightly, perhaps due to individual risk aversion. A simple Chi-square test was employed to test whether the proportions of subjects submitting bids of $50 per share or greater differed across the monetary and non-monetary experiments. Based on this test, one could not the reject the null hypothesis of equal proportions at the 5% significance level.

As in the case of acquirer bids, no statistical distinction can be made between target offers across the monetary and non-monetary experiments. A comparison of Figures 3.3a and 4.3b shows closely similar distributions of offers for the six listed target values. The sole qualitative difference between the distributions is that under monetary payoffs the frequency of "aggressive" offers in excess of the company's acquisition value, $P_T > w(v)$,

is reduced (about 20%, down from 33%). In Figure 4.3b, the average offers for the values v=0,...v=100 are 18.3, 37.7, 58.0, 81.2, 104.9, 125.1 respectively. These average about $5 per share lower than offers under non-monetary payoffs. Again, this may be due to risk aversion on the part of subjects. As before, the offers conform to the hypothesis of naive bargaining behavior -- target offers in the interval [v, w(v)]. Note also that most target offers fall well short of the empirically best (i.e. expected profit-maximizing) response given the distribution of actual acquirer acceptence prices.

Discussion. The experimental results indicate that both acquirers and targets systematically deviate from normative behavior. These deviations, however, are far more serious in the case of the uninformed party (the acquirer). A sizeable minority of acquirers in Version 1, 50% of acquirers in Version 2, and the vast majority of acquirers in Version 3 made bids which resulted in expected losses.[7] In each case, a significant number of subjects employed a naive bargaining strategy -- bidding a price equal to E[v], the unconditional value of the company under target management. Indeed, comparing the three versions, the incidence of the winner's curse appears to depend directly on the difficulty of the inferential task at hand. Loss-making prices are relatively easy to identify and avoid in Version 1, less so in Version 2. Finally, the acquirer's 50% comparative advantage for the company, disguises the fact that all price offers are loss-making in Version 3.

When given the opportunity to name a sale price, targets choose price demands that are more "cooperative" than would be predicted under a normative (equilibrium) model of behavior. In contrast to the equilibrium prediction, the majority of price demands fall short of the potential acquisition value of the firm, thereby providing potential profits to the acquirer. Indeed, this

cooperative demand behavior is not a best response against the actual distribution of acquirer acceptence prices. As demonstrated, targets can significantly increase their expected profits by employing more aggressive demand strategies (i.e., demanding higher prices given the true value of the company).

The adoption of cooperative pricing strategies by targets, though individually suboptimal, is to the advantage of the parties collectively. Given the pattern of target pricing behavior, acquirers can earn positive profits by naming high acceptence prices. In fact, one can determine that the acquirer's empirically best acceptence prices are \$110, \$125, and \$130 in Versions 1, 2, and 3 respectively.[8] For these acceptence prices, agreements are virtually assured. Thus, the experimental evidence indicates that target price setting is more successful than acquirer price setting in generating mutually beneficial transactions. (Under acquirer offers, in the majority of agreements across the versions, the acquirer falls prey to the winner's curse.) Thus, the experiments provide some evidence in favor of the common practice of price setting by informed sellers.

5. Concluding Remarks.

The key feature in the bargaining setting is the presence of asymmetric information -- the fact that the target has better information about the value of the company than does the acquirer. The experimental evidence suggests that many acquirers failed to adjust their bids/offers accordingly in light of this difference in information. As a result, a sizable minority of subjects in Version 1 and the vast majority in Version 2 suffered expected losses in the exercise. Once recognized, the asymmetry of information should put the less well-informed party on its guard. For instance, the management board of

the acquiring firm might undertake the following line of reasoning in making an offer:

> If we offer a price that is readily accepted, what inference can we draw about the target's true worth? Are we going to be sold a "bill of goods"? After all, it is more likely that target management will unload an ailing company than a healthy one. If our bid is accepted, will we be getting what we paid for or even less than what we paid for?

Though we are all familiar with the advice "Buyer Beware", our results indicate that individuals fail to carry out the inferences necessary to apply formally this maxim in bargaining under uncertainty.

One might expect that the incidence of the winner's curse in actual negotiation settings would be mitigated by a process of learning from one's mistakes. For instance, a past record of acquisitions which proved less profitable than expected ought to make a firm more cautious in the future. This observation suggests an important extension to the "one shot" experiments described herein. In sequential experiments, each subject would face a series of acquisition decisions and would be informed of the result and payoff of a current decision before undertaking a subsequent one. It is natural to ask whether, given this feedback, acquirers would conform more closely to optimal bid behavior (or at least avoid loss-making transactions). A similar series employing target offers would be useful for addressing the issue of convergence to equilibrium behavior and the possible role of seller reputations over time. Sequential experiments such as these would provide a powerful test concerning the persistence of suboptimal behavior given the opportunity for subject learning.

The experiments suggest that a common setting for the winner's curse (perhaps the most common) is buyer-seller bargaining. Value uncertainty is the norm in a variety of settings ranging from corporate acquisitions to the settlement of law suits. While the bargaining process holds the promise of

mutually beneficial agreements, under the winner's curse these profits may prove illusory to one or even both parties. For instance, with the increase in the number of corporate takeovers in the 1980's has come evidence that acquiring companies have often paid too much for what they're getting. As many as one-third of all acquisitions prove to be failures and an additional one-third fail to live up to expectations (Wall Street Journal, 1981). Our analysis suggests that potential acquirers should temper their optimism by recognizing that the most profitable targets will be the ones that refuse to sell out. Consequently, companies actually acquired will be worth less than first thought. Indeed, the greatest danger of falling prey to the winner's curse occurs during a bidding war over a takeover candidate. Some of the largest and most publicized takeover contests -- Dupont's winning bid for Conoco against the competition of Mobil and Seagram, for instance — have resulted in purchase premiums more than 50% above market prices. The success or failure of these acquisitions remains to be seen.

Notes

* We would like to thank Vernon Smith and especially an anonymous referee for constructive comments on this paper.

1. This view was first espoused by Ronald Coase (1960). Hoffman and Spitzer provide extensive references on the Coase "theorem".

2. As another example, Perrow (1984) has provided substantial evidence that the primary reason for ship accidents concerns the naive assumptions each ship makes about the likely behaviors of the opposing ship. Neale and Bazerman (1983) have also shown that negotiators fail to seek out the information available from considering the perspective of the other party in negotiation, and this lack of perspective taking ability hinders the success of the focal negotiator.

3. Version 3 is an adaptation to the bargaining setting of Akerlof's (1970) well-known example of adverse selection in a market for "lemons".

4. First-and-final offers were employed for two reasons. First, their very simplicity insures that subjects easily grasp the experiment. Second, Samuelson (in press) shows that of all bargaining procedures, a first-and-final offer maximizes the offerer's expected profit in equilibrium. In short, this is an optimal strategy for both the informed and uninformed party.

5. The offer P = \$30 generates a return of \$0 with probability .7 and provides positive profits uniformly distributed between \$30 and \$0 with probability .3. By comparison, the offer P = \$20 generates a return of \$0 with probability .8 and gives profits uniformly distributed between \$20 and \$0 with probability .2.Finally, the offer P = \$40 generates a return of \$0 with probability .6 and profits uniformly distributed between \$30 and -\$10 with probability .4. Thus, it follows that a \$30 offer stochastically dominates a \$20 or a \$40 offer (and, by the same reasoning, any other alternative).

6. For a proof that this is the seller's optimal bargaining procedure, see Samuelson (in press).

7. To test subject understanding, the work packet contained a series of questions to be answered after a price offer had been chosen. In each version, subjects were constrained to choose a preferred price offer from among the trio of \$70, \$50, or \$0 per share. Presented this less complicated choice, subjects showed a slight improvement in performance. For instance, in Version 1, 25% of subjects made a loss-making offer, \$50 per share. In Version 3, thirteen percent of subjects offered the correct normative bid, a zero price.

A second series of short questions placed subjects in the role of target and asked whther they'd accept or reject specific acquirer offers. All subjects adopted the target's dominant strategy — accepting the acquirer's demand if and only if P_T exceeded v.

8. As Figures 3.1b, 3.2b, and 3.3b, acquirer acceptence prices are widely spread in the respective versions. Furthermore, given the pattern of target price demands, these acceptence prices imply varying profit consequences. In Version 1, acceptence prices below \$65 per share are loss-making. Prices in this lower range were submitted by nearly 50% of subjects, who fell prey to the winner's curse. In Version 2, all acceptence prices are profitable. Seemingly subjects perceived this, since nearly three-quarters of acceptence prices exceeded \$80 per share. Finally, in Version 3, acceptence prices below \$80 per share are loss-making. As under acquirer price offers, a majority of subjects continue to fall prey to the winner's curse in Version 3.

References

Akerlof, G., "The Market for Lemons: Qualitative Uncertainty and the Market Mechanism," Quarterly Journal of Economics, 89, August 1970, 488-500.

Bazerman, M.H., "Negotiator Judgment: A Critical Look at the Rationality Assumption," American Behavioral Scientist, 27, Nov.-Dec. 1983, 211-228.

Bazerman, M.H., & Samuelson, W.F., "I Won the Auction But Don't Want the Prize," Journal of Conflict Resolution, 27, December 1983, 618-634.

Capen, E.C., R.V. Clapp, and W.M. Campbell, "Competitive Bidding in High Risk Situations," Journal of Petroleum Technology, 23, June 1971, 641-53.

Coase, R., "The Problem of Social Cost," Journal of Law and Economics, 3, October 1960, 1-44.

Hoffman, E. and M.L. Spitzer, "The Coase Theorem: Some Experimental Evidence," Journal of Law and Economics, 25, April 1982, 73-98.

Miller, E.M., "Risk, Uncertainty, and Divergence of Opinion," Journal of Finance, 32, September 1977, 1151-68.

Neale, M.A., & Bazerman, M.H., "The Effect of Perspective Taking Ability on the Negotiation Process Under Different Forms of Arbitration," Industrial and Labor Relations Review, 36, April 1983, 378-388.

Perrow, C., Normal Accidents, New York: Free Press, 1984.

Raiffa, H., The Art and Science of Negotiation, Cambridge: Harvard University Press, 1982.

Samuelson, W., "Bargaining Under Asymmetric Information," Econometrica, in press.

Shubik, M., "The Dollar Auction Game: A Paradox in Non-Cooperative Behavior and Escalation," Journal of Conflict Resolution, 15, March 1971, 109-111.

Smith, V. "Microeconomic Systems as Experimental Science," American Economic Review, 72, December 1982, 923-955.

Teger, A.I., Too Much Invested to Quit, New York: Pergamon, 1979.

Wall Street Journal, "To Win a Bidding War Doesn't Insure Success of Merged Companies," September 1, 1981, p.1.

Appendix.

Acquiring a Company
(Version 1)

In the following exercise, you will represent Company A (the acquirer) which is currently considering acquiring Company T (the target) by means of a tender offer. You plan to tender in cash for 100% of Company T's shares but are unsure how high a price to offer. The main complication is this: the value of the company depends directly on the outcome of a major oil exploration project it is currently undertaking.

If there is an oil strike the company will be worth $60/share under current management or $130/share in Company A's hands. However, if there is no oil strike, the company will be worth only $30/share whether it is managed by Company A or Company T. You have had two oil exploration firms give you there best estimates of the probability of an oil strike. Both firms have told you that there is a 50% chance of a strike.

Value of Stock

	Under Company A	Under Company T
Oil Strike	$130/share	$60/share
No Oil Strike	$30/share	$30/share

Note, if an oil strike occurs, the potential value of the company under Company A's direction is much greater than its value under Company T.

The board of directors of Company A has asked you to determine the price they should offer for Company T's shares. This offer must be made now, before the outcome of the drilling project (strike or no strike) is known. From all indications, Company T would be happy to be acquired by Company A, provided it is at a profitable price. You expect Company T to delay a decision on your bid until the results of the project are in, then accept or reject your offer before releasing the news of the drilling results to the press.

Thus, you (Company A) will not know the results of the exploration project when submitting your price offer, but Company T will know the results when deciding whether or not to accept your offer. In addition, Company T is expected to accept any offer by Company A that is greater or equal to the (per share) value of the company under its own management.

As the representative of Company A, you are deliberating over price offers in the range $0/share (this is tantamount to making no offer at all) to $150/share. What price offer per share would you tender for Company T's stock?

My Tender Price is: $_____ per share.

Do not turn the page until you respond above.

Now suppose that you represent the target company and it has the opportunity to name its price (which the acquiring company can accept or reject). All other aspects of the situation are as described earlier. Now you are privy to Company T's inside informaation, i.e. you know the company's actual value per share. For each of the share values listed below, specify the share price you would demand on behalf of company T. If your demand is accepted by company A, you (Company T) will earn a profit equal to the excess of the price you demanded over the company's current value.

If Company T's per Share value is:	My price demand is:
$30	_____
$60	_____

Finally, suppose once again that you represent the acquiring company. You are deliberating whther or not to accept the price quoted by the target. What is the maximum price named by the target that you would be willing to accept? (An answer of $x per share means that you will accept any price that is smaller or equal to $x and will reject any price greater than x.)

As Company A's representative, the highest price I would be willing to accept is:

$_____ per share.

Acquiring a Company
(Version 2)

This version is identical to Version 1 with two exceptions. Version 2 replaces the second paragraph of Version 1 (including the table of values) with the following paragraph.

> The very viability of Company T depends on the exploration outcome. In the worst case (if the exploration fails completely), the company under current management will be worth nothing -- $0/share. In the best case (a complete success), the value under current management could be as high as $100/share. Given the range of exploration outcomes, all share values between $0 and $100 per share are considered equally likely. By all estimates the company will be worth considerably more in the hands of Company A than under current management. In fact, whatever the value under current management. the company will be worth $30 per share more under the management of Company A than under Company T. In the worst case, the company is worth $0/share under either management. If the exploration project generates a $50/share value under Company T, the value under Company A is $80/share. Similarly, a $100/share value under Company T implies a $130/share value under Company A, and so on.

The second difference is that as representative of the target company, the subject is asked to name price demands for the values, 0, 20, 40, 60, 80, and 100.

Acquiring a Company
(Version 3)

This version is identical to Version 1 with two exceptions. Version 3 replaces the second paragraph of Version 1 (including the table of values) with the following paragraph.

> The very viability of Company T depends on the exploration outcome. In the worst case (if the exploration fails completely), the company under current management will be worth nothing -- $0/share. In the best case (a complete success), the value under current management could be as high as $100/share. Given the range of exploration outcomes, all share values between $0 and $100 per share are considered equally likely. By all estimates the company will be worth considerably more in the hands of Company A than under current management. In fact, whatever the value under current management. the company will be worth 50% more under the management of Company A than under Company T. In the worst case, the company is worth $0/share under either management. If the exploration project generates a $50/share value under Company T, the value under Company A is $75/share. Similarly, a $100/share value under Company T implies a $150/share value under Company A, and so on.

The second difference is that as representative of the target company, the subject is asked to name price demands for the values, 0, 20, 40, 60, 80, and 100.

Acquiring a Company
(Version 3 - Monetary Payoffs)

This version is identical to the prior Version 3 with two exceptions. First, since the exercise involved potential monetary gains (and possibly) losses, subjects indicated first whether they chose to participate. The exercise began with the following cover sheet:

(Name and/or ID Number)

Acquiring a Company

The following exercise (beginning on page 2) has been used for classroom purposes with over 400 B.U. MBA students. Should you choose to participate in today's version, you will be faced with choices that involve actual gains and (possibly) losses of small sums of money. The exercise takes no more than 5 to 10 minutes to complete.

If, at the present moment, you know categorically that you do not wish to participate in the exercise, check the box below. If you choose to participate, it is still your decision as to what monetary risks to take. At the conclusion of the exercise, the experimenters will pay you any winnings immediately. If you lose, they will accept your payment or an IOU.

☐ I DO NOT WISH TO PARTICIPATE.

If you check the box above, feel free to glance through the packet, but please do not write on any of the pages.

Second, in the monetary experiments, subjects were assigned to represent either the acquiring company or the target company. Thus, the second page of the handout described the negotiation situation (exactly the same information as before) from the viewpoint of the company which the subject was representing. A final page explained how monetary payoffs would be determined. For the role of acquirer, the text read as follows:

Monetary Payoffs. After you make your offer and all papers are collected, the experimenters will determine your monetary payoff in the following manner:

1. The value of the company under target management will be determined by a random draw from a jar containing all integer values between $0 and $100 per share, inclusive. It will then be determined whether your offer is accepted by the target company (as described above).

2. If your offer is not accepted, your score and monetary payoff are $0.

3. If your offer is accepted:

a) Your score is the difference between the computed value of the acquired company under your management and the price you offered.

b) Your monetary payoff is 10¢ for each point that you score.

EXAMPLE: You make an offer of $50 per share which (based on the random draw) is accepted with the result that you acquire a company worth $60 per share to you. Your score is 60 - 50 = 10. Your monetary receipt is 10¢.

If you were to acquire a company worth less to you than your offer price, your score would be negative, and you would owe the experimenters 10¢ for each negative point.

Please fill in the following item below.

My Tender Price is: $_____ per share.

For the role of the target, the text read as follows:

Monetary Payoffs. Below, you will be asked to report the offer you would make for each of the values $0, $20, $40, $60, $80, $100 per share. After you write down your offers and all papers are collected, the experimenters will determine your monetary payoff in the following manner:

1. The value of the company under target management will be determined by a random draw from a jar containing the integer values $0, $20, $40, $60, $80 and $100 per share. Only your written offer for the single value drawn will be pertinent to the calculation of your payoff.

2. The maximum price acceptable to Company A will be drawn from a second jar. The jar will contain the actual maximum acceptence prices submitted by other MBA subjects in an earlier experiment using this same exercise.

3. If Company A's maximum acceptable price is greater or equal to your price offer, your offer is accepted. Otherwise it is rejected.

   a) If your offer is rejected, your score and monetary payoff are $0.

   b) If your offer is accepted:

   Your score is the difference between your price offer and the value of the company under your management.

   Your monetary payoff is 10¢ for each point that you score.

   EXAMPLE: The value $60 is drawn for which you named a price of $70 per share. If Company A's maximum acceptable price (drawn from a second jar) is smaller than $70, your offer is rejected and your payoff is $0. If Company A's maximum price is greater than $70, your offer of $70 is accepted and your profit score is 70 - 60 = 10. Your monetary receipt is (10)(10¢) or $1.00.

Now please fill in the following items below.

The oil results indicate Company A's current value to be:	My Price Offer is:
$0 per share	$_____ per share.
$20 per share	$_____ per share.
$40 per share	$_____ per share.
$60 per share	$_____ per share.
$80 per share	$_____ per share.
$100 per share	$_____ per share.

Figure 1.1 The Distribution of Acquirer Bids in Version 1.

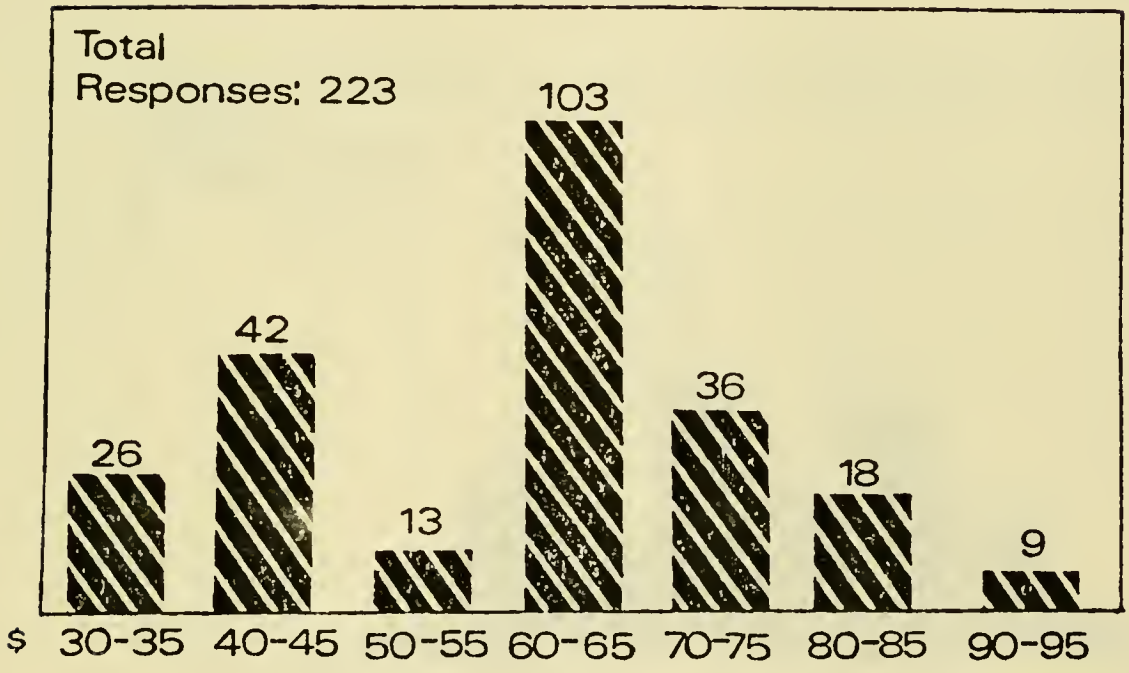

Figure 1.2 The Distribution of Acquirer Bids in Version 2.

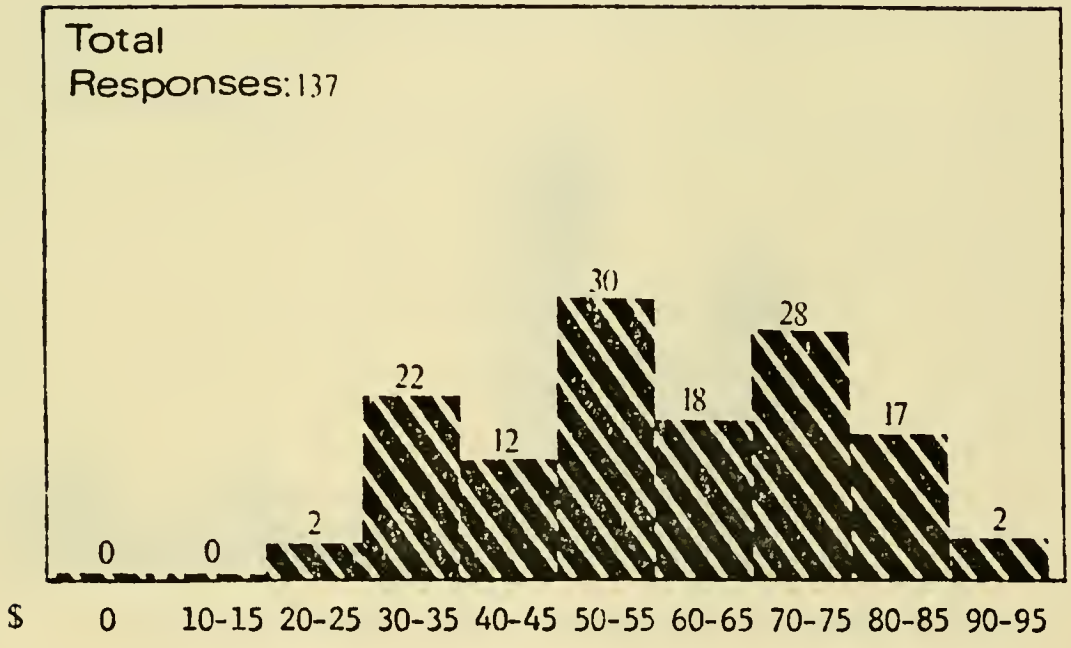

Figure 1.3 The Distribution of Acquirer Bids in Version 3.

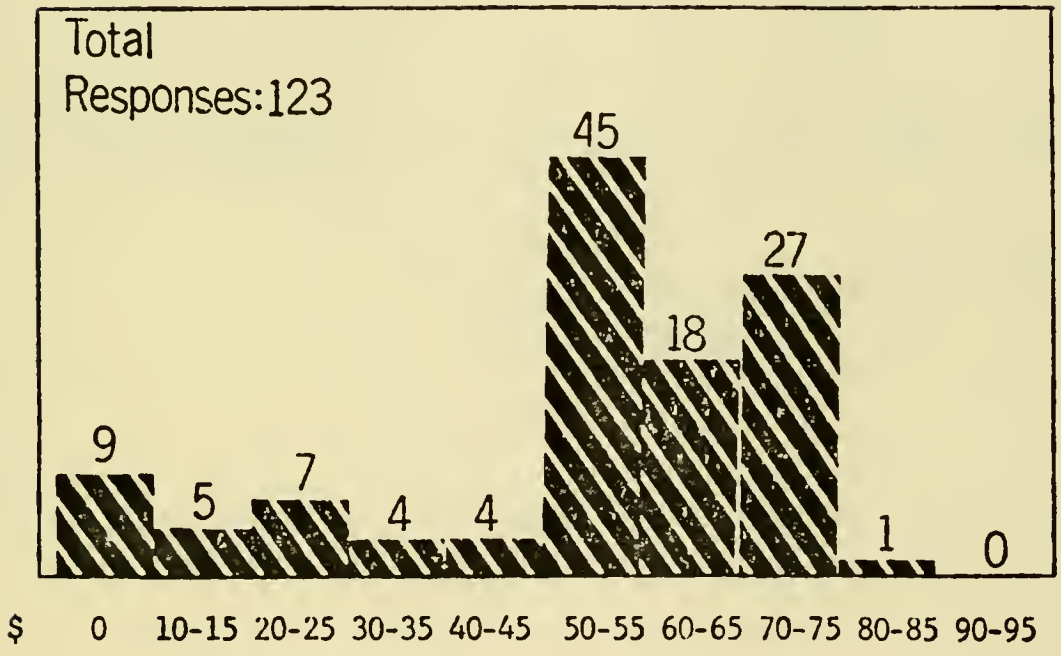

Figure 2.1 The Acquirer's Expected Profit in Version 1.

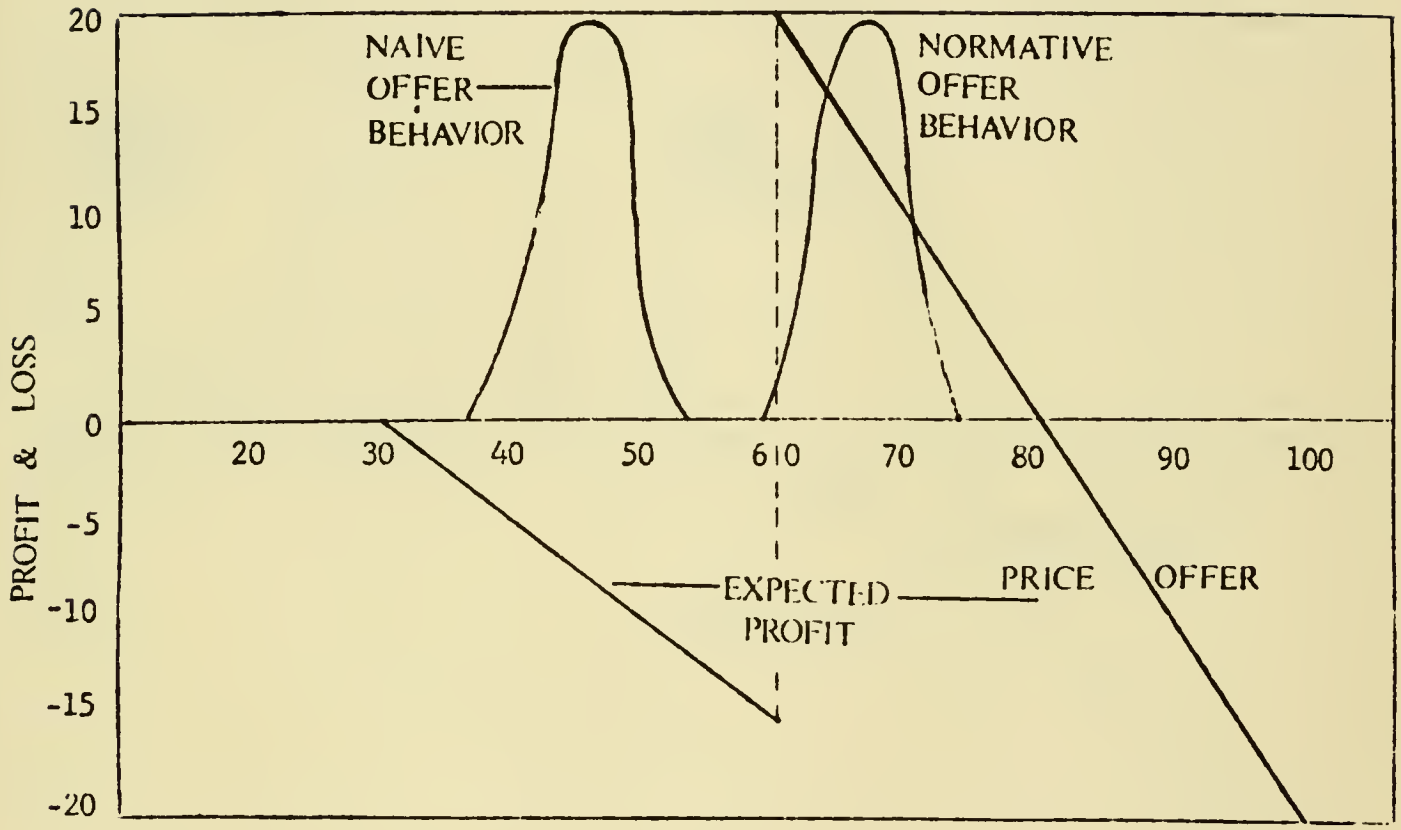

Figure 2.2 The Acquirer's Expected Profit in Version 2.

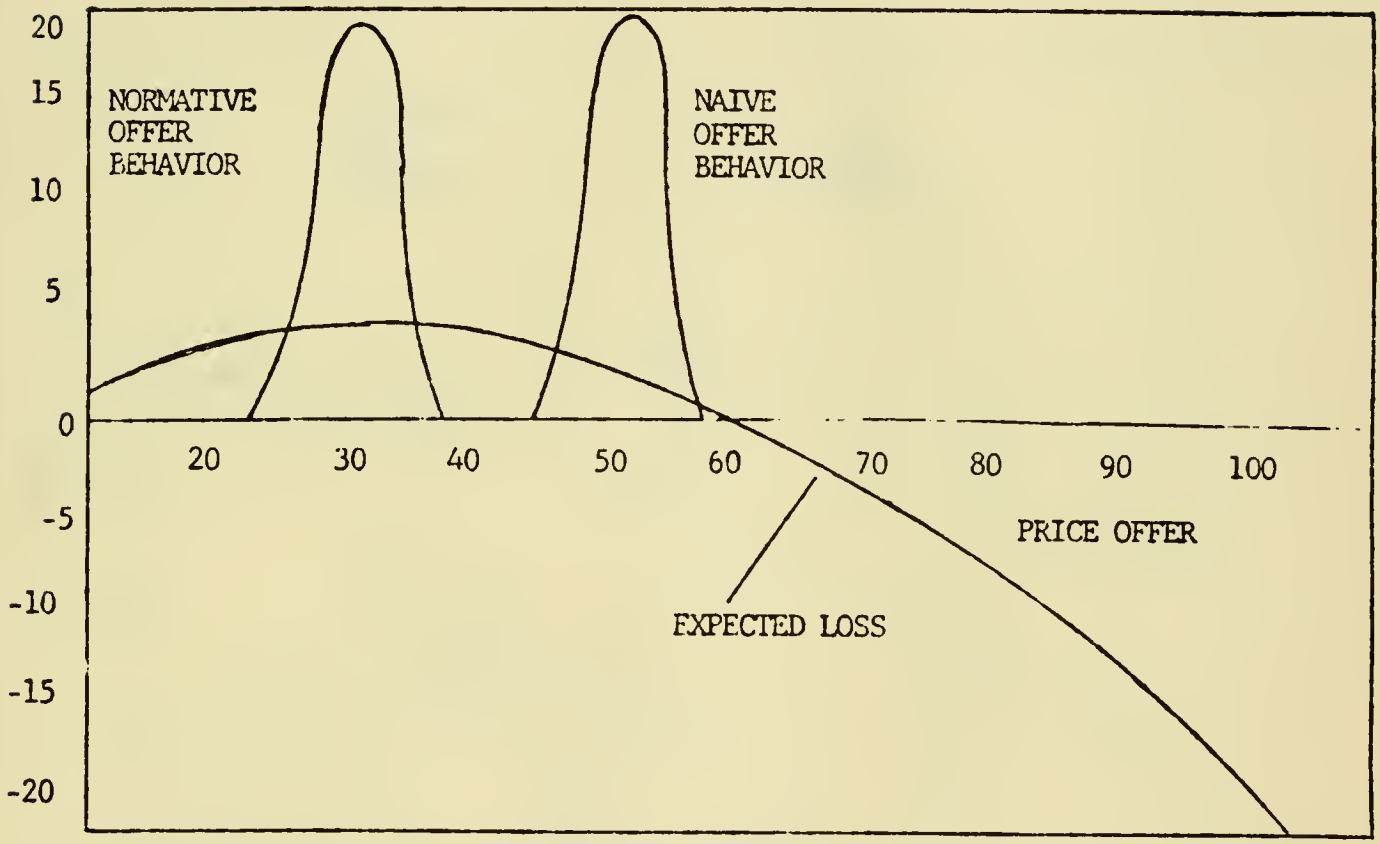

Figure 2.3 The Acquirer's Expected Profit in Version 3.

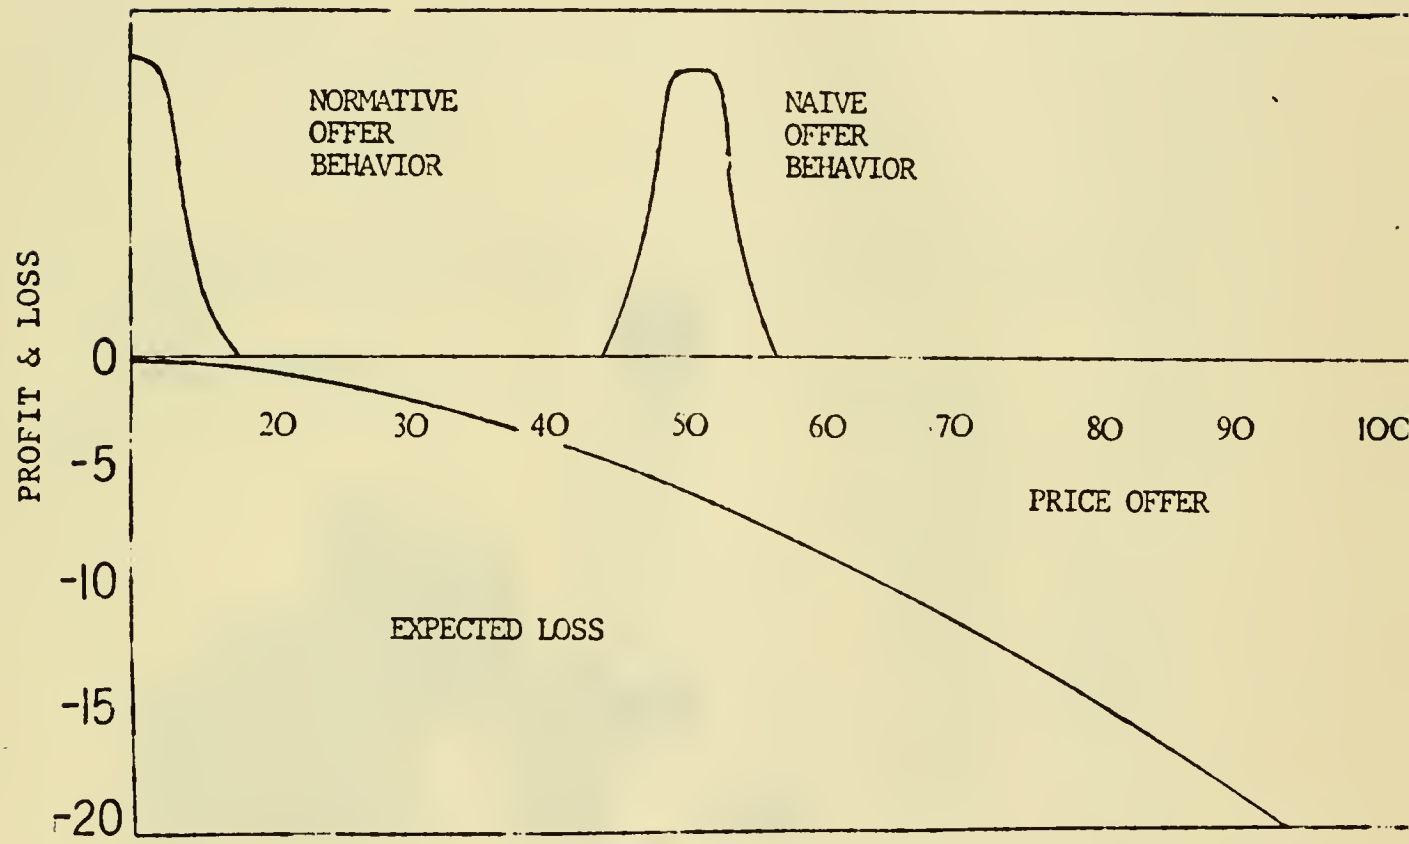

Figure 3.1a The Distribution of Target Offers in Version 1.

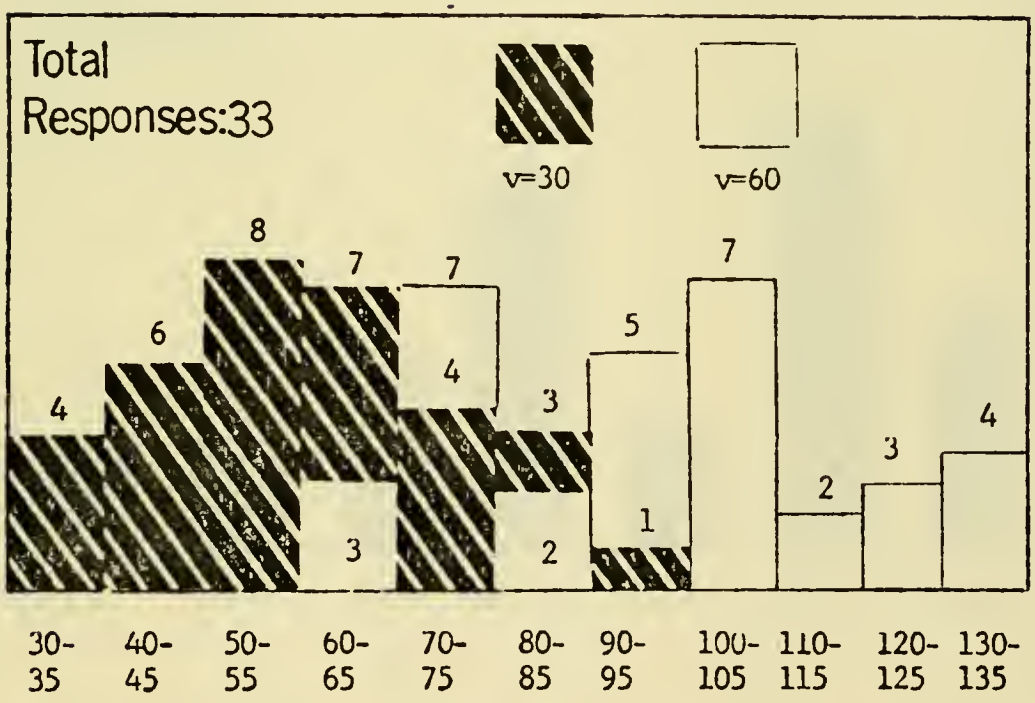

Figure 3.1b The Distribution of Acquirer Acceptence Prices in Version 1.

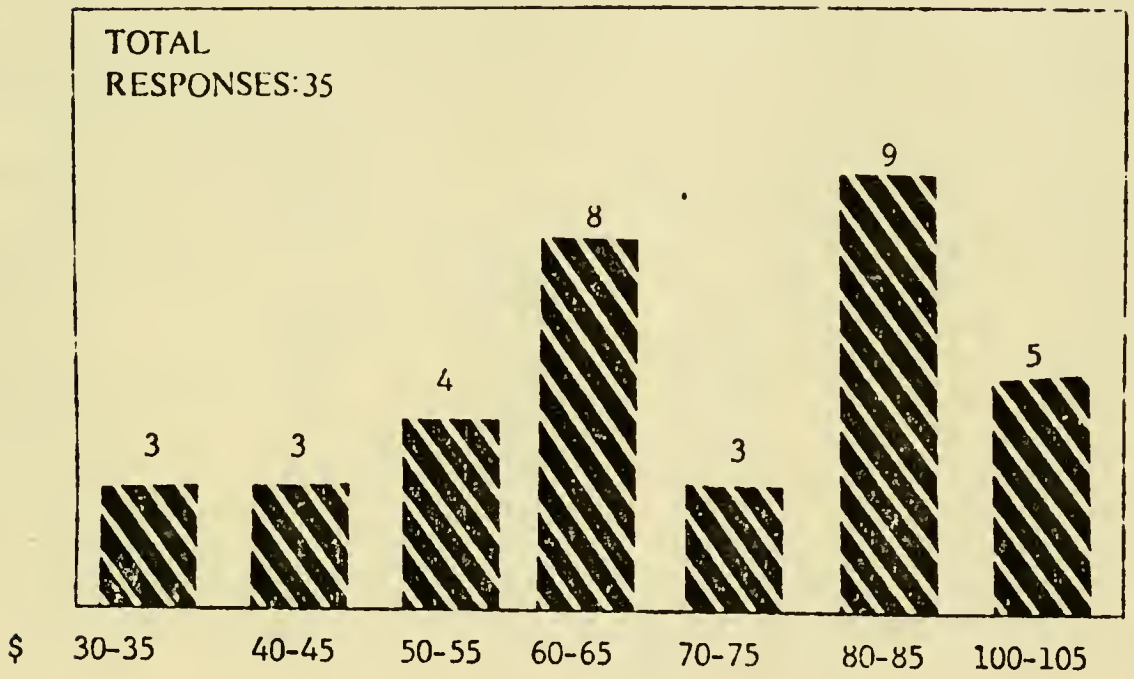

Figure 3.2a The Distribution of Target Offers in Version 2.

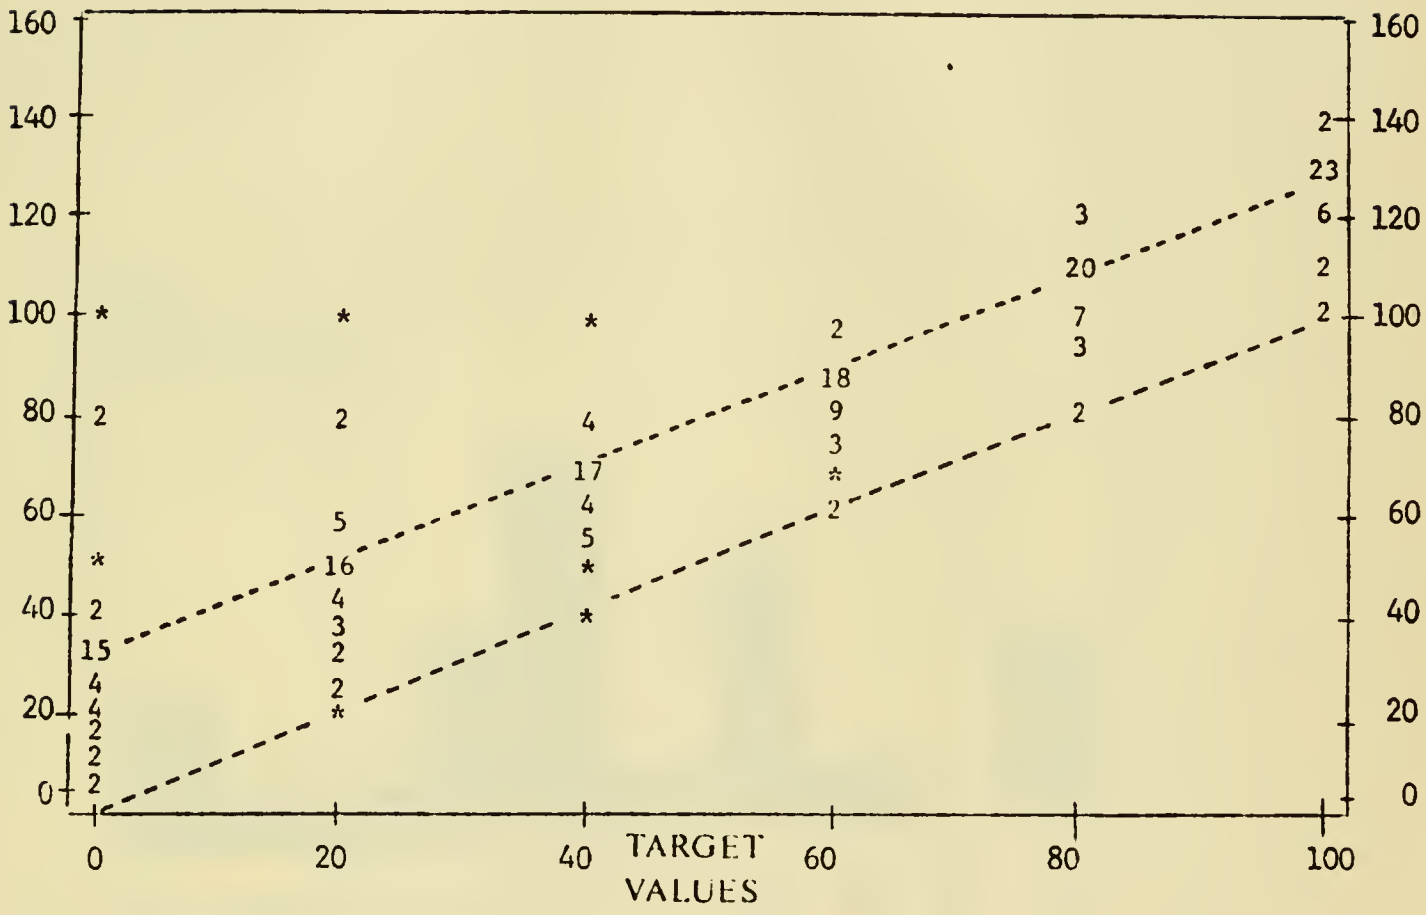

Figure 3.2b The Distribution of Acquirer Acceptence Prices in Version 2.

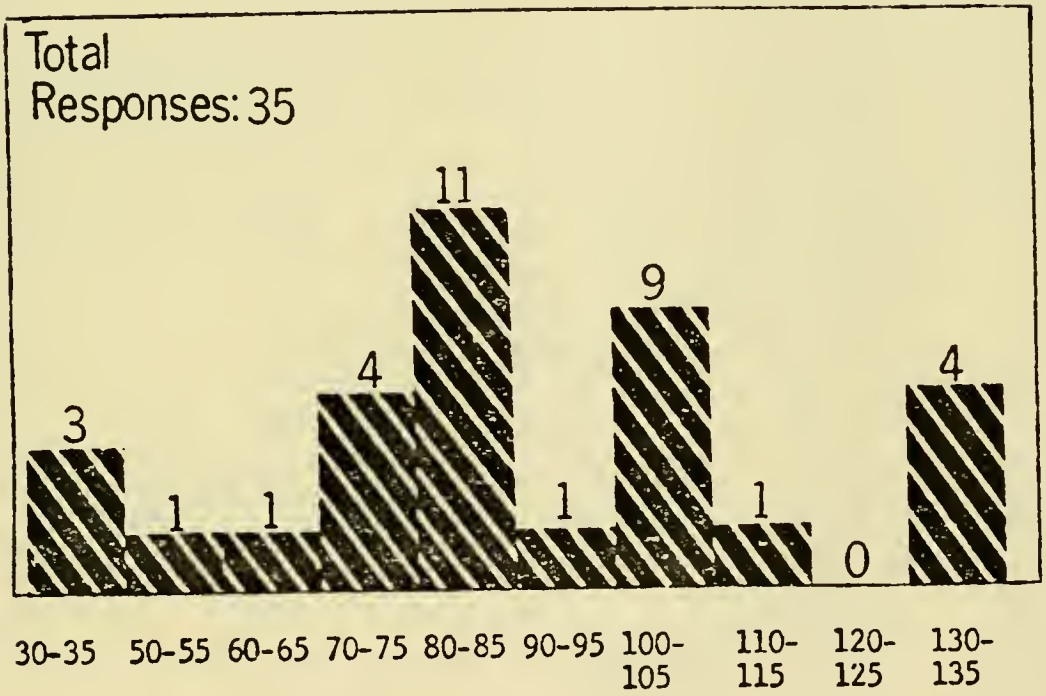

Figure 3.3a The Distribution of Target Offers in Version 3.

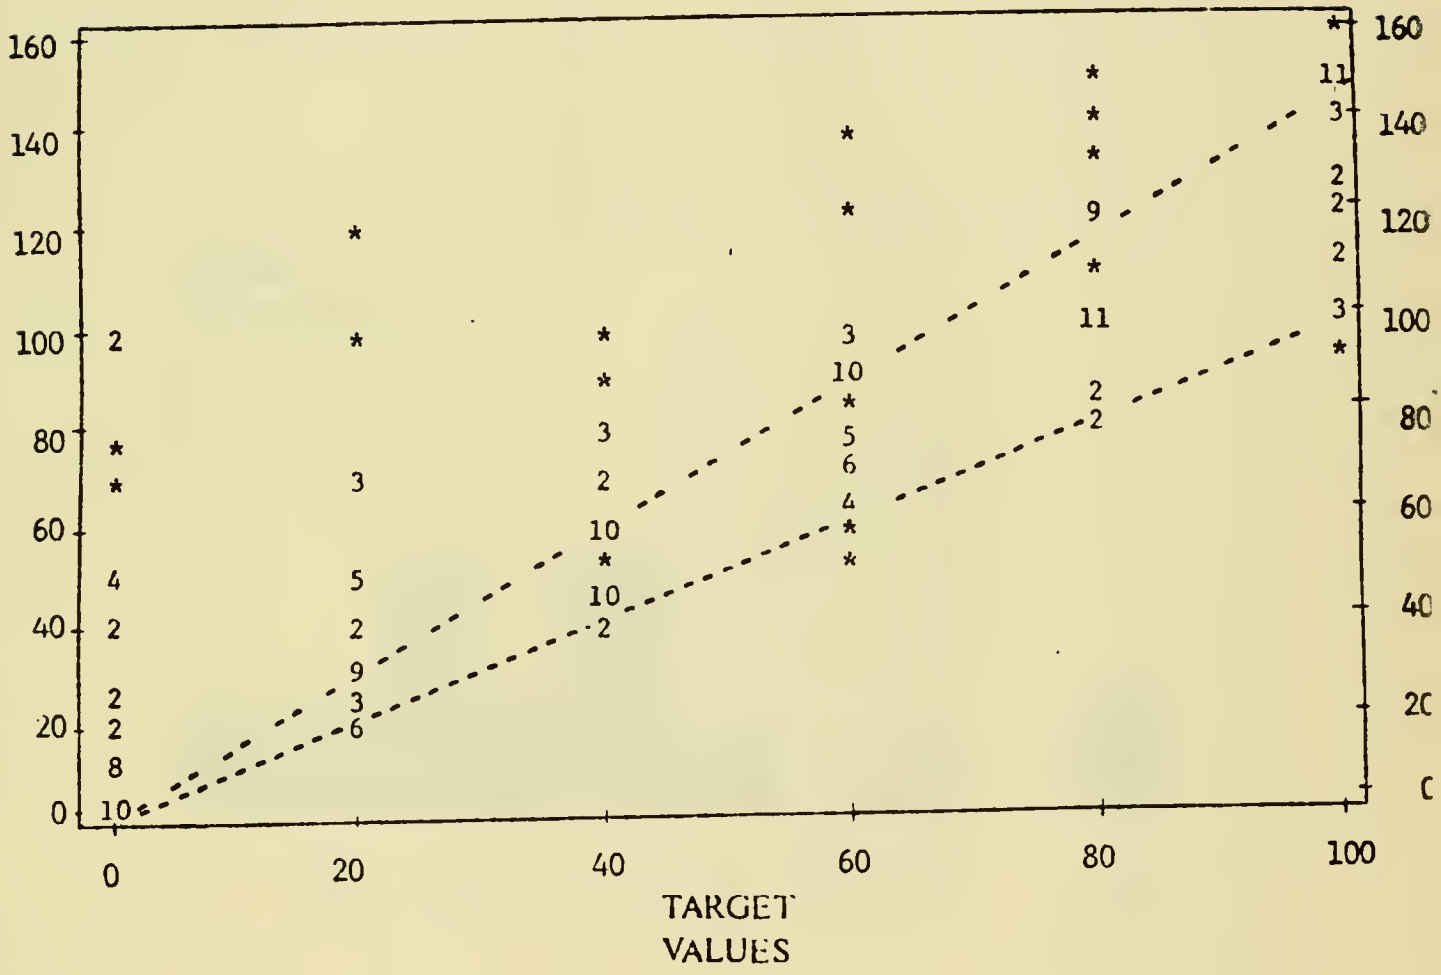

Figure 3.3b The Distribution of Acquirer Acceptence Prices in Version 3.

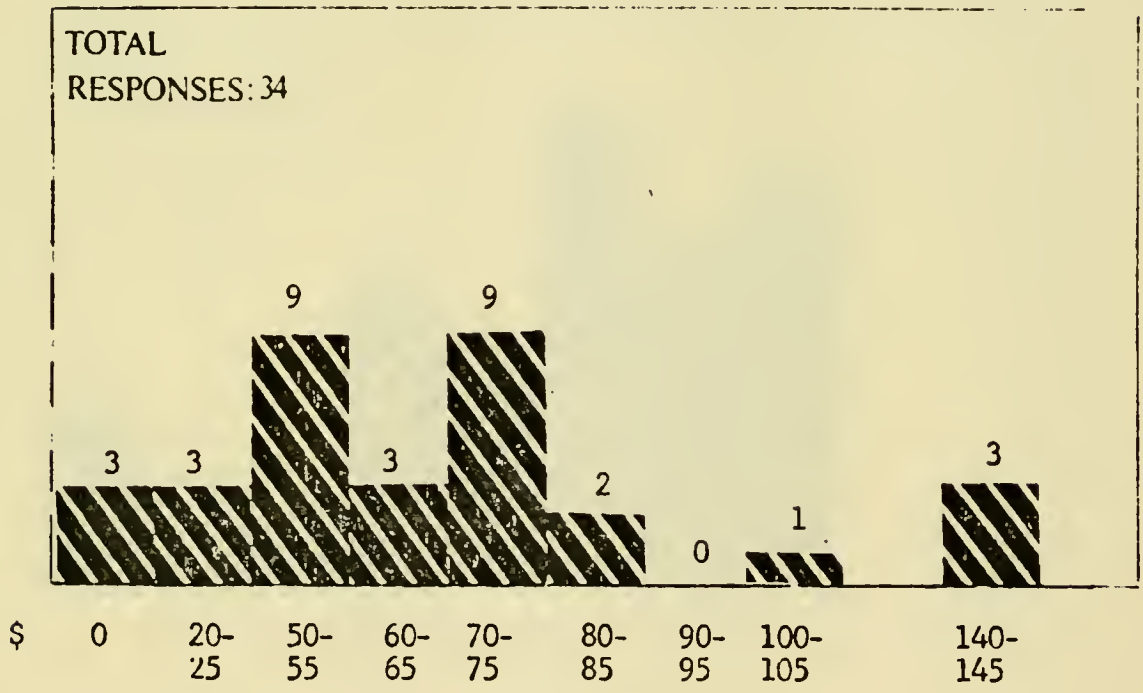

Figure 4.3a Acquirer Bids in Version 3 with Monetary Payoffs.

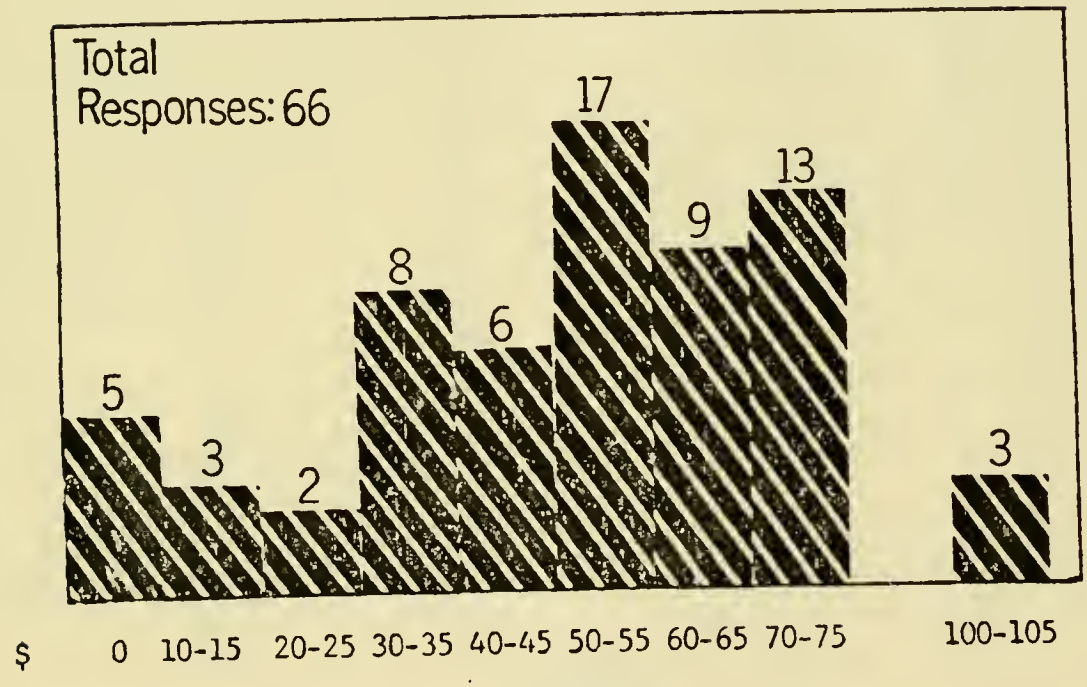

Figure 4.3b Target Offers in Version 3 with Monetary Payoffs.

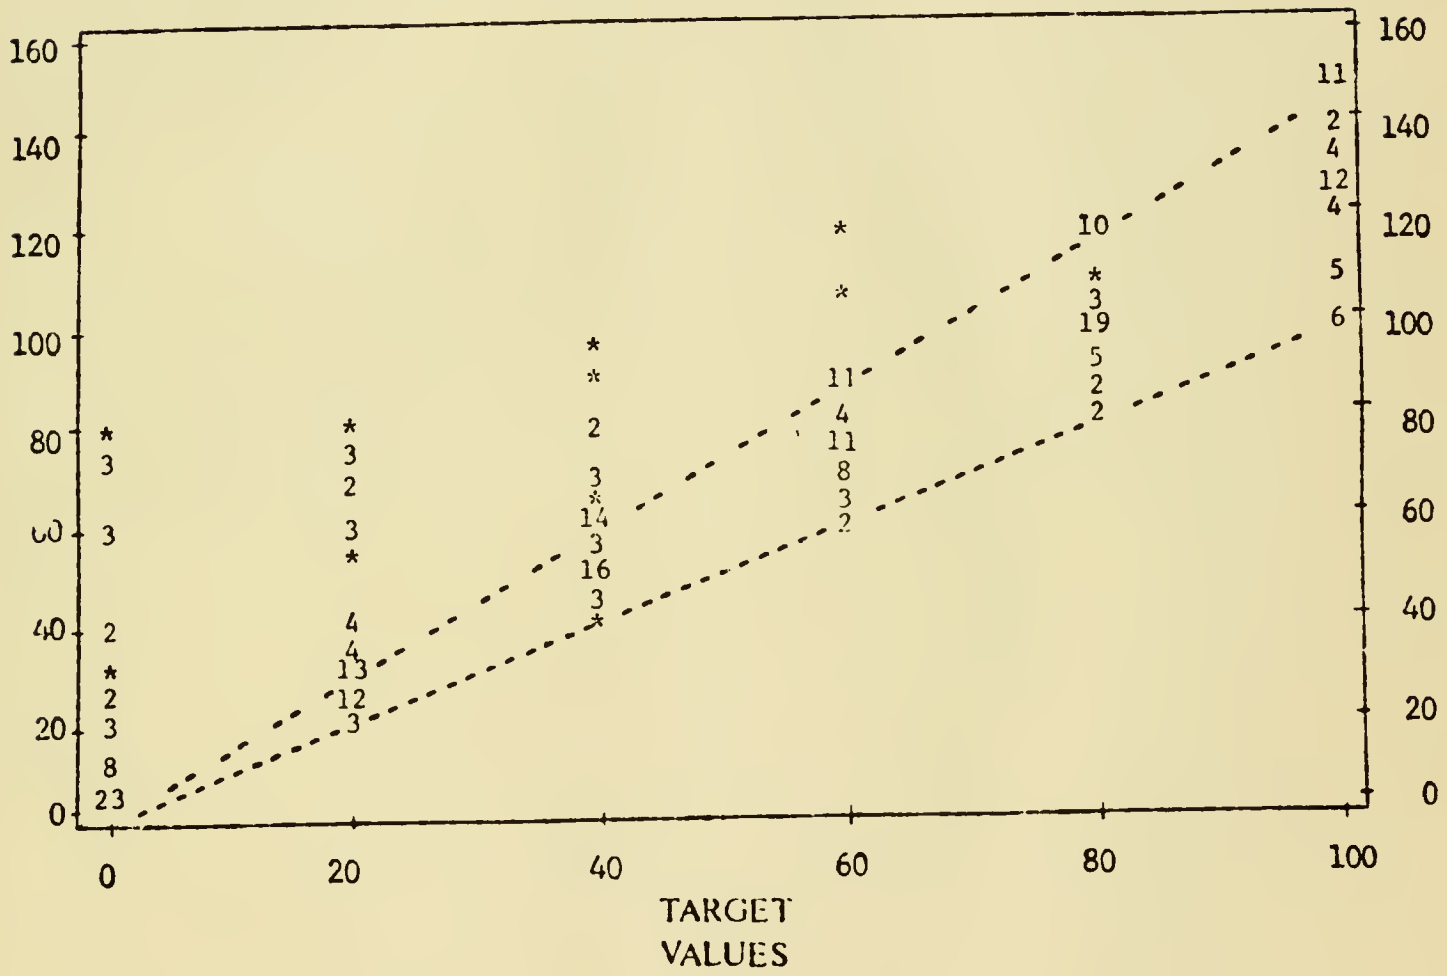

4110 070

MIT LIBRARIES
3 9080 004 480 783

Date Due

SEC

NOV 16 1989

Lib-26-67

BASEMENT

BASEMENT

Bar Code is
on Last
Page